

# An Incremental Learner for Language-Based Anomaly Detection in XML

Harald Lampesberger

Department of Secure Information Systems  
University of Applied Sciences Upper Austria  
Email: harald.lampesberger@fh-hagenberg.at

**Abstract**—The Extensible Markup Language (XML) is a complex language, and consequently, XML-based protocols are susceptible to entire classes of implicit and explicit security problems. Message formats in XML-based protocols are usually specified in XML Schema, and as a first-line defense, schema validation should reject malformed input. However, extension points in most protocol specifications break validation. Extension points are wildcards and considered best practice for loose composition, but they also enable an attacker to add unchecked content in a document, e.g., for a signature wrapping attack.

This paper introduces datatyped XML visibly pushdown automata (dXVPAs) as language representation for mixed-content XML and presents an incremental learner that infers a dXVPA from example documents. The learner generalizes XML types and datatypes in terms of automaton states and transitions, and an inferred dXVPA converges to a good-enough approximation of the true language. The automaton is free from extension points and capable of stream validation, e.g., as an anomaly detector for XML-based protocols. For dealing with adversarial training data, two scenarios of poisoning are considered: a poisoning attack is either uncovered at a later time or remains hidden. Unlearning can therefore remove an identified poisoning attack from a dXVPA, and sanitization trims low-frequent states and transitions to get rid of hidden attacks. All algorithms have been evaluated in four scenarios, including a web service implemented in Apache Axis2 and Apache Rampart, where attacks have been simulated. In all scenarios, the learned automaton had zero false positives and outperformed traditional schema validation.

**Keywords**—XML, grammatical inference, visibly pushdown automata, stream validation, anomaly detection, experimental evaluation.

## I. INTRODUCTION

The Extensible Markup Language (XML) [1] is ubiquitous in electronic communication, e.g., the Simple Object Access Protocol (SOAP), the Extensible Messaging and Presence Protocol (XMPP), the Security Assertion Markup Language (SAML), and many data serialization formats. The success of XML boils down to its rich data models and tool support: Instead of specifying some protocol from scratch, a software developer can simply define a subset of XML and reuse existing parsing and querying tools.

XML attacks, in particular, the signature wrapping attack [2], have motivated this work. The signature wrapping attack exploits identity constraints and best practices for composition in XML Schema (XSD) [3], and the attack's goal is to modify a document without violating a cryptographic signature. Several high-value targets were vulnerable

over the years, e.g., the Amazon EC2 cloud control SOAP interface [4] and many SAML frameworks [5]. Attack tool support is already available [6], and fixing the problem tends to be hard [7]–[10]. Signature wrapping is a showcase for language-theoretic security because it is the result of design choices. A document with references is logically not a tree but often wrongly treated as such in modular software, and the need for determinism in composed schemas has led to extension points in many specifications as attack enablers.

This paper extends a previous grammatical inference approach, where a language representation is learned from example documents [11]. Use cases for the presented approach are anomaly detection in XML-based protocols and schema inference for interface hardening. The contributions are automaton models as language representations for mixed-content XML, algorithms for datatype inference from texts, an incremental learner, and an experimental evaluation.

For representing event streams of mixed-content XML, the proposed datatyped XML visibly pushdown automata (dXVPAs) and character-data XVPAs (cXVPAs) introduce transitions for text contents in the original XVPA definition by Kumar et al. [12]. The proposed learner converges to a good-enough language approximation in terms of a dXVPA. An inferred automaton for stream validation mitigates the signature wrapping attack because it is free from extension points. Counting the mind changes between incremental steps is a heuristic for measuring convergence in the learning progress. Furthermore, the learner has been designed with poisoning attacks in mind. Two scenarios are considered: a successful poisoning attack is uncovered at some later time and a poisoning attack stays hidden but is statistically rare [13]. For the first case, the learner provides an unlearning operation, and for the second case, a sanitization operation trims low-frequent states and transitions from an automaton. All algorithms have been implemented and evaluated in four scenarios, where various XML attacks are simulated: two synthetic and two realistic scenarios utilizing Apache Axis2 and Rampart. In all scenarios, the learned automaton outperformed baseline schema validation.

## A. XML

XML specifies a syntax: open- and close-tags for elements, attributes, namespaces, allowed characters for text content and attribute values, processing instructions for the

parser, inline Document Type Definitions (DTD) [1], and comments. The syntax allows ambiguities, e.g., an element without text content, and XML Information Set [14] therefore defines a data model to remove syntactic ambiguities: A document has an infoset if it is *well formed* and all namespace constraints are satisfied.

Business logic accesses infoset items in a document through an interface. Common APIs for XML can be distinguished into (a) stream based, e.g., Simple API for XML (SAX) and Streaming API for XML (StAX) [15] and (b) tree based, e.g., a Document Object Tree (DOM).

A *schema* is basically a grammar, and the XML community provides several schema languages for specifying production rules, e.g., DTD, XSD, Relax NG [16], and Extended DTD (EDTD) [12] as a generalization. Productions are of form  $a \rightarrow B$ , where  $B$  is a regular expression and called *content model* of  $a$ . In DTD, rules are expressed over elements. To raise expressiveness, productions in XSD, Relax NG, and EDTD are defined over *types*, where every type maps to an element. This mapping is surjective: two types can map to the same element.

*Schema validation* is checking language acceptance of a document. *Typing* is stricter than validation by assigning a *unique type* from productions to every element [17]. The power of regular expressions and the surjective relation between types and elements can introduce ambiguity and nondeterminism, but determinism is desired, e.g., for assigning semantics to types. DTD and XSD therefore have syntactic restrictions to ensure deterministic typing. Schema validation and typing are first-line defenses against attacks; however, XML identity constraints and extension points in XSD can render validation ineffective.

### B. Language-Theoretic Vulnerabilities

The XML syntax is context free and infoset items are tree structured, but a document is not always logically a tree. Identity constraints like keys (ID) and key references (IDREF, IDREFS) introduce self-references that go beyond context freeness. Cyclic and sequential references turn a finite tree data model logically into an infinite tree, and operations such as queries become computationally harder [18]. Furthermore, XSD introduces additional constraints (unique, key, and keyref) over text contents, attribute values, and combinations thereof. Checking identity constraints during schema validation is costly because indices need to be constructed, or the data model is traversed many times.

Also, there are two philosophies of modularity in XSD: *schema subtyping* [19] by refining productions and *schema extension points* using wildcards (`xs:any`). Extension points allow loose coupling and are considered best practice [20]. In an XSD, a wildcard is often accompanied by the `processContents="lax"` attribute which has a tremendous effect on validation: if there is no schema in the parser's search space for a qualified element at an extension point,

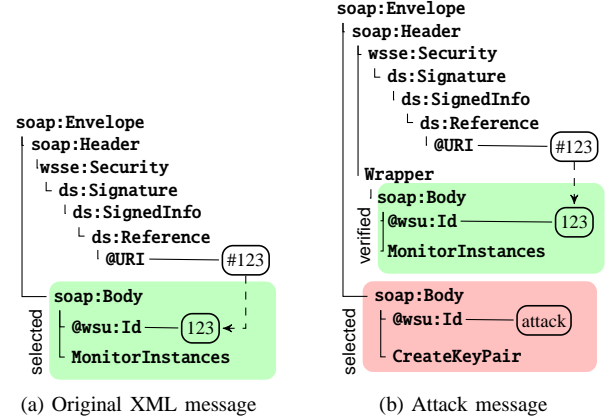


Figure 1. XML signature wrapping attack (reproduced from [4])

validation is skipped. By choosing a random namespace, an attacker can place arbitrary content at an extension point. Unfortunately, extension points are present in many standards, e.g., SOAP, XMPP, and SAML.

### C. XML Attacks

Attacks can be distinguished into *parsing* and *semantic* attacks. Parsing attacks target lexical and syntactical analyses, e.g., for Denial-of-Service. Examples are oversized tokens, high node counts from unbounded repetitions [21], and coercive parsing [22]. Schema validation is unable to reject parsing attacks when they are placed at an extension point. A special class of parsing attacks originates from inline DOCTYPE declarations, i.e., exponential entity expansion, external entities for privilege escalation, and server-side request forgery (SSRF) [23].

Semantic attacks aim for misinterpretation, e.g., by tampering with structure and texts. CDATA fields [22] can exclude reserved characters (e.g. angled brackets) from lexical analysis as a helper for many semantic attacks, e.g., XML, SQL, LDAP, XPath, and command injection; path traversal; memory corruption in interacting components; and cross-site request forgery (XSRF) and cross-site scripting (XSS) with respect to web applications [21].

*1) Signature Wrapping Attack:* Signature wrapping is a semantic attack. XML Signature [24] specifies a `ds:signature` element that holds one or more hashes of referenced resources (i.e., elements in the document) and is signed for authenticity. The resources are referenced by an ID or an XPath expression. Signature checking verifies the authenticity of the `ds:signature` element and compares the stored hashes with computed ones. Checking is usually treated as a Boolean decision, independent from the business logic, and a vulnerability emerges when verified document locations are not communicated between software modules accessing the document.

In a signature wrapping attack, the referenced resource is moved to an extension point, and a malicious element is

placed instead at the original location. An example based on the Amazon EC2 attack [4] is shown in Figure 1. As a precondition, the attacker needs access to some signed document (Figure 1a). The SOAP schema has an extension point in `soap:Header`, and the original message body is hidden in a wrapper element for skipping schema validation. The signature remains valid (Figure 1b), and the business logic wrongly processes the attacker-provided message body.

2) *Signature Wrapping Countermeasures: Security policies* [2] can enforce properties of SOAP messages, but policy checking is computationally costly. Gruschka and Iacono [25] furthermore show a successful signature wrapping attack on Amazon EC2 that satisfies security policies.

Rahaman et al. [26] propose an *inline approach* by adding an element to SOAP headers that stores document characteristics. Unfortunately, if a single element in the header is not signed, the approach can also be circumvented [7].

Gajek et al. [7] and Somorovsky et al. [5] propose *improved signature verification* by returning a filtered document view, but the business logic needs to be adapted accordingly. Gajek et al. [8] also propose FastXPath for location-aware XPath-based references in signatures. Namespace injection in XPath-based references could eventually break this countermeasure too [9].

Jensen et al. [10] propose *schema hardening* by removing extension points and restricting repetitions. Hardening is effective because elements cannot be hidden anymore; however, all composed schemas need to be known beforehand, and generating a single unified hardened schema is computationally hard. Experiments have also shown a significant slowdown in schema validation.

#### D. Research Questions

Removal of extension points is an effective countermeasure, but compiling a unified schema is difficult [10]. This paper therefore proposes a monitor for an XML-based system. The monitor has a learner and validator component. The learner component infers a dXVPA, and the validator component utilizes an optimized variant of the automaton to validate documents sent to the system under observation. Validation is relative to the training data, and the approach is therefore called *language-based anomaly detection*. If the validator component rejects a document, some filtering or extended policy checking could be performed, but these operations are not in the scope of this work.

The assumed attacker is capable of reading and modifying documents in transit and sending a malicious document directly to the system under observation.

Clients and services are considered black boxes, where message semantics are unknown to the monitor; however, semantics are important under the language-theoretic security threat model because an attack is basically a misinterpretation. When a system interprets a document, semantics for elements and texts are derived from assigned types

and datatypes respectively, where types and datatypes are usually defined in software (ad hoc) or in schema production rules. Attacks affect at least one type or datatype in a document for causing misinterpretation. The system under observation is assumed to have *type-consistent behavior*: for all manifestations of an expected type or datatype in a document, the behavior is well specified. In other words, language-based anomaly detection only works if attacks are syntactically distinguishable from expected types and datatypes. To sum up, the research questions are:

- RQ1 What is a suitable language representation for types and datatypes that is capable of stream validation?
- RQ2 Can this language representation be learned?
- RQ3 Can the proposed approach identify attacks?

#### E. Methodology

1) *Language Representation*: In mixed-content XML, texts are strings over Unicode, and they are allowed between a start- and an end-tag, an end- and a start-tag, two start-tags, and two end-tags. XSD provides datatypes for specifying texts, where every datatype has a value space and a lexical space over Unicode. A language representation that captures document structure and texts needs to be expressive with respect to typing and support *stream validation* for open-ended XML protocols (i.e., XMPP) and very large documents. To answer RQ1, the paper introduces dXVPAs as an extension of XVPAs [12]. XVPAs are known to recognize StAX event streams for linear-time stream validation, but text contents are not considered yet. A dXVPA introduces transitions for datatypes of text content, and a cXVPA is an optimized dXVPA representation for linear-time stream validation in the validator component.

2) *Learning from Positive Examples*: The learner component receives examples and computes automata for the validation component. This learning setting corresponds to Gold's *identification in the limit from positive examples* [27], and according to Fernau [28], the definition is as follows.

**Definition 1** (Identification in the limit from positive examples [28]). Let  $\mathcal{L}$  be a target language class that can be characterized by a class of language-describing devices  $\mathcal{D}$ .  $E: \mathbb{N} \rightarrow L$  is an enumeration of strings for a language  $L \in \mathcal{L}$ , and the examples may be in arbitrary order with possible repetitions. Target class  $\mathcal{L}$  is *identifiable in the limit* if there exists an inductive inference machine or learner  $I$ :

- Learner  $I$  receives examples  $E(1), E(2), \dots$
- Learner  $I$  reacts by computing a stream of hypotheses (e.g., automata)  $D_1, D_2, \dots$  such that  $D_i \in \mathcal{D}$ .
- For every enumeration of  $L \in \mathcal{L}$ , there is a convergence point  $N(E)$  such that  $L = L(D_{N(E)})$  and  $j \geq N(E) \implies D_j = D_{N(E)}$ .

RQ2 is answered by specifying algorithms for inferring datatypes from text and automata from documents. Further-




Figure 2. Incremental learning step

more, unlearning and sanitization operations for dealing with adversarial training data are provided.

3) *Experimental Evaluation*: A learning-based approach is still a heuristic and requires experimental evaluation. Four datasets have been generated: two synthetic ones using a stochastic XML generator and two realistic ones from a web service implemented in Apache Axis2 and Apache Rampart. The service has been implemented according to best practices, and attacks have been performed manually and automatically by the WS-Attacker [6] tool. Detection performance, learning progress in terms of mind changes, and the effects of unlearning and sanitization have been analyzed to answer RQ3.

## II. GRAMMATICAL INFERENCE OF XML

Figure 2 illustrates the incremental learning step. The learner component maintains an internal visibly pushdown automaton (VPA)  $A$  and counters  $\omega_\delta, \omega_Q, \omega_F$  for transitions, states, and final states. A VPA is a special pushdown automaton with three disjoint alphabets: a call alphabet that pushes on the stack, an internal alphabet that leaves the stack unchanged, and a return alphabet that pops from the stack. This concept originates from program analysis, and for XML, the alphabets represent different kinds of events. The set of states is implicitly the stack alphabet. For a complete definition of VPAs, the reader is directed to Alur et al. [29].

Algorithm 3 (*incWeightedVPA*) receives a document event stream  $E(i)$  and updates the VPA and the counters. The counters are frequencies of states and transitions from training data and necessary for unlearning and sanitization operations. Algorithm 4 (*trim*) removes zero-weight states and transitions, and Algorithm 5 (*genXVPA*) constructs a minimized dXVPA. The dXVPA becomes an optimized cXVPA for the validator component, and acceptance of documents can then be efficiently decided.

### A. Document Event Stream

**Definition 2** (Document event stream). A *document event stream*  $w$  is a sequence of StAX events  $e$  carrying values  $lab(e)$ . There are three kinds of events: *startElement* and *endElement* for open- and close-tags of qualified element names and *characters* for texts. Processing instructions, comments, and entity references are ignored. Attributes are

alphabetically sorted, treated as elements with a leading @ symbol, and mapped to a subsequence of *startElement*, *characters*, and *endElement* events.

For simpler notation, a *startElement* event for qualified element  $m$  is denoted as  $m$ , and  $\bar{m}$  is the respective *endElement* event. The value of a *characters* event is a string over Unicode, and nested CDATA sections are automatically unwrapped by the parser.

XSD provides datatypes for specifying text contents. In this work, only the lexical spaces of XSD datatypes [30] are considered in a generalized notation of lexical datatypes.

**Definition 3** (Lexical datatypes). Let  $T$  be a set of *lexical datatypes*. A lexical space is a regular language over Unicode  $U$ , and  $\phi : T \rightarrow REG(U)$  assigns lexical spaces.

Lexical datatypes allow to define datatyped event streams, where datatypes replace text contents in *characters* events.

**Definition 4** (Datatyped event stream). A *datatyped event stream*  $w'$  is a sequence of *startElement*, *endElement*, and *characters* events. The value of a *characters* event  $e$  is a datatype  $lab(e) \in T$ . A document event stream  $w$  corresponds to a datatyped event stream  $w'$  if  $w$  and  $w'$  have congruent event kinds, the qualified element names in *startElement* and *endElement* events are the same, and text content in a *characters* event in  $w$  is in the lexical space of the congruent *characters* event in  $w'$ .

### B. Language Representation

1) *Datatyped XVPA*: The dXVPAs extend XVPAs [12] with datatypes, so they can accept datatyped event streams.

**Definition 5** (dXVPA). A dXVPA  $A$  over  $(\Sigma, M, \mu, T, \phi)$  is a tuple  $A = (\{Q_m, e_m, X_m, \delta_m\}_{m \in M}, m_0, F)$ .  $\Sigma$  is a set of qualified element names,  $M$  is a set of modules (equivalent to types in schemas),  $\mu : M \rightarrow \Sigma$  is a surjective mapping that assigns elements to modules,  $T$  is a set of datatypes, and  $\phi : T \rightarrow REG(U)$  assigns lexical spaces over Unicode.

For every module  $m \in M$ :

- $Q_m$  is a finite set of module states
- $e_m \in Q_m$  is the module's single entry state
- $X_m \subseteq Q_m$  are the module's exit states
- $\delta_m = \delta_m^{call} \uplus \delta_m^{int} \uplus \delta_m^{ret}$  are module transitions
  - $\delta_m^{call} \subseteq \{q_m \xrightarrow{c/q_m} e_n \mid n \in \mu^{-1}(c)\}$ , where  $c$  is a *startElement* event value that pushes  $q_m$  on the stack
  - $\delta_m^{int} \subseteq \{q_m \xrightarrow{\tau} p_m \mid \tau \in T\}$  and  $\tau$  is the value of a datatyped *characters* event
  - $\delta_m^{ret} \subseteq \{q_m \xrightarrow{\bar{c}/p_n} q_n \mid n \in \mu^{-1}(c)\}$ , where  $\bar{c}$  is an *endElement* event value that pops  $p_n$  from the stack; the relation is deterministic, i.e.,  $q_n = q'_n$  whenever  $q_m \xrightarrow{\bar{c}/p_n} q_n$  and  $q_m \xrightarrow{\bar{c}/p_n} q'_n$



The minimally required datatypes are computed by Algorithm 1 (*minLex*). The algorithm terminates after  $|T|$  steps in the worst case. Acceptance of a string by a datatype is checked in topological sort order with respect to  $\leq_{lex}$ . To minimize the number of checks, a candidates set *cand* is maintained. If  $w$  is in some lexical space,  $w$  is also in all greater datatypes because  $\leq_{lex}$  is transitive, and the up-set can be removed from *cand*. Furthermore, the topological order guarantees that the matched datatypes are minimal and incomparable. Algorithm *minLex* always returns a nonempty set because the  $\top$  datatype has space  $U^*$  and matches for

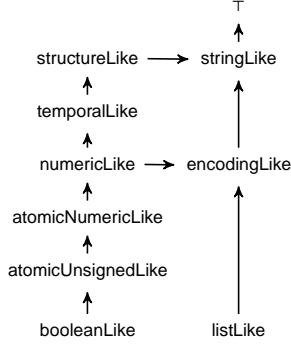


Figure 5. Ordering  $\leq_s$  on kinds of lexical datatypes

any string.

2) *Preference Heuristic*: Figure 4 already suggests that lexical spaces of XSD datatypes are often incomparable and ambiguous. This leads to weird datatype choices, e.g.,  $\minLex(\text{false}) = \{\text{language, boolean, NCName}\}$ . The antichain is lexically correct, but some datatypes are semantically more informative and preferred over others. A second step in datatype inference is therefore to drop the least informative datatypes from minimally required datatypes. The proposed heuristic captures the XSD type hierarchy and datatype semantics in an ordering  $\leq_s$  for kinds of datatypes. Figure 5 illustrates the ordering, and kinds are defined as:

```

stringLike = {string, normalizedString, token, ENTITY, ID, IDREF, NMTOKEN}
listLike = {ENTITIES, IDREFS, NMTOKENS}
structureLike = {anyURI, NOTATION, QName, Name, language, NCName}
encodingLike = {base64Binary, hexBinary}
temporalLike = {gDay, gMonth, gYear, gYearMonth, gMonthDay, date, duration,
  time, dayTimeDuration, yearMonthDuration, dateTime, dateTimeStamp}
numericLike = {nonPositiveInteger, nonNegativeInteger, positiveInteger,
  decimal, integer, negativeInteger}
atomicNumericLike = {float, double, long, int, short, byte}
atomicUnsignedLike = {unsignedLong, unsignedInt, unsignedShort,
  unsignedByte}
booleanLike = {boolean}

```

There is also a distinguished  $\top$  kind for the  $\top$  datatype for upward closure. Algorithm 2 (*pref*) compares pairs of minimally required datatypes, and if two datatypes are comparable with respect to  $\leq_s$ , the greater datatype is removed from the set. The resulting set  $R'$  is still an antichain of datatypes with respect to  $\leq_{lex}$ .

3) *Datatypes Event Stream for Learning*: For learning, every text in a document event stream needs to be mapped to its minimally required datatypes:

$$\minReq(w) = \text{pref}(\minLex(w)) \quad \text{for string } w \quad (1)$$

$$\text{dtyped}(e) = \begin{cases} \minReq(\text{lab}(e)) & \text{if characters} \\ e & \text{for other events} \end{cases} \quad (2)$$

---

#### Algorithm 2: *pref*

---

**Input:** lexical datatype system  $(T, \phi, \sim_s, \leq_s)$   
 datatypes  $R \subseteq T$

**Output:** preferred datatypes  $R' \subseteq T$

---

```

1  $R' := R$ 
2 for  $\tau, \tau'$  in  $R$  and  $\tau \neq \tau'$  do
3   if  $[\tau]_{\sim_s} <_s [\tau']_{\sim_s}$  then  $R' := R' \setminus \{\tau'\}$ 

```

---

The learner also needs to be able to aggregate minimally required datatypes from different text contents. Let  $v, w$  be to strings over Unicode. The minimally required datatypes that accept both strings are:

$$\minReq(v, w) = \max_{\leq_{lex}} \minReq(v) \cup \minReq(w) \quad (3)$$

The  $\max_{\leq_{lex}}$  operation guarantees a nonempty antichain with respect to  $\leq_{lex}$  that cover both strings.

**Example 2.** Let  $S = \{1, \emptyset, \text{true}, 33\}$  be Unicode strings, then  $\minReq(S) = \{\text{boolean, unsignedByte}\}$ .

#### D. The Incremental Learner

A famous result by Gold [27] states that the language class of unrestricted regular expressions is not learnable in the limit from positive examples only. This result translates to dXVPAs because modules characterize regular languages over types and datatypes. The full language class of datatyped event streams expressible in dXVPAs can therefore not be learned from example documents only, and restrictions are necessary. Two restrictions originating from schema complexity are considered:

- *Simplicity of regular expressions.* Bex et al. [31] have examined 202 DTDs and XSDs and conclude that the majority of regular expressions in practical schema productions are simple because types occur only a small number of times in expressions.
- *Locality of typing contexts.* Martens et al. [17] have studied 819 DTDs and XSDs from the web and XML standards, and typing elements in 98% is local, i.e., the type of an element only depends on its parent.

To capture simplicity, Bex et al. [32] define the class of single-occurrence regular expressions (SOREs). In a SORE, a symbol occurs at most once, and the majority of schema productions in the wild belong to this class. SOREs generate a 2-testable regular language, and  $k$ -testable regular languages [33] are known to be efficiently learnable from positive examples only.

A  $k$ -testable regular language is fully characterized by a finite set of allowed substrings of length  $k$ , and learning is collecting the substrings. This can be done efficiently by constructing a prefix tree acceptor (PTA), i.e., a DFA that accepts exactly the examples, and *naming the states* according to the string prefixes that lead to them. Merging




Figure 6. Typing of an element

states whose names share the same  $(k-1)$ -length suffix then generalizes the automaton. This can be done implicitly while constructing a PTA. The proposed learner utilizes this idea by embedding typing information in state names that are derived from prefixes of datatyped event streams.

1) *Typing Mechanisms*: Typing can be thought of as a function that determines the type of an element from its name and other elements in the document [17], [34].

Figure 6 illustrates typing mechanisms by representing the info-set of a document without text contents and identity constraints as a tree, where  $lab(v)$  is the qualified element name of node  $v$ . Efficient stream processing requires deterministic typing, and Martens et al. [17] therefore define *1-pass preorder typing* (1PPT): a schema allows 1PPT if the type of every node  $v$  can be determined from the *preceding*( $v$ ) subtree as shown in Figure 6a. The authors surprisingly show that typing based on the ancestor-sibling string *anc-lsib-str*( $v$ ) is sufficient for the 1PPT property.

Let  $lsib(v) = lab(u_1) \cdots lab(u_m) \cdot lab(v)$  be a left-sibling string, where  $u_1, \dots, u_m$  are the left siblings of  $v$ . The ancestor-sibling string is then  $anc-lsib-str(v) = lsib(i_1) \# lsib(i_2) \# \cdots \# lsib(i_n)$  such that  $i_1$  is the root node,  $i_n = v$ , and  $i_{j+1}$  is a child of  $i_j$ . An example is in Figure 6b.

Element Declaration Consistency (EDC) and Unique Particle Attribution (UPA) are syntactic restrictions for productions in XSD to ensure deterministic typing. These restrictions are tighter than necessary for the 1PPT property. In XSD, a node  $v$  is typed by the ancestor string *anc-str*( $v$ ) as shown in Figure 6c. The ancestor string is defined as  $anc-str(v) = lab(i_1) \cdot lab(i_2) \cdots lab(i_n)$ , where  $i_1$  is the root node,  $i_n = v$ , and  $i_{j+1}$  is a child of  $i_j$ .

2) *Incremental Update*: Named states in Algorithm 3 (*incWeightedVPA*) are the foundation for state merging. The algorithm iterates over document event stream  $w$  in a single pass and returns an updated VPA and counters. In a run, the algorithm maintains a stack, collects element names, and for every event, a next state is derived from three *state naming functions* with signatures  $call : Q \times \Sigma \rightarrow Q$ ,  $int : Q \rightarrow Q$ , and  $ret : Q \times Q \times \Sigma \rightarrow Q$ . A transition is then stored to connect the current with the next state.

The three functions utilize the discussed typing mechanisms, and two *state naming schemes* are proposed.

**Definition 10** (State naming schemes). A state is a pair  $(u, v)$

### Algorithm 3: *incWeightedVPA*

---

**Input:** VPA  $A = (Q, q_0, F, Q, \delta)$  over  $\Sigma \uplus T \uplus \bar{\Sigma}$   
lexical datatype system  $(T, \phi, \sim_s, \leq_s)$   
state naming functions  $call, int, ret$   
counters  $\omega_Q, \omega_F, \omega_\delta$   
document event stream  $w$

**Output:** updated VPA  $A$  and counters  $\omega_Q, \omega_F, \omega_\delta$

---

```

1   $s := \perp$  // empty stack
2   $q := q_0$  // current state
3  for  $e$  in  $dtyped(w)$  do
4      switch  $eventType(e)$  do
5          case startElement do
6               $\Sigma := \Sigma \cup \{lab(e)\}$ 
7               $q' := call(q, lab(e))$ 
8               $\omega_Q(q') := \omega_Q(q') + 1$ 
9               $\delta^{call} := \delta^{call} \cup \{q \xrightarrow{lab(e)/q} q'\}$ 
10              $\omega_\delta(q \xrightarrow{lab(e)/q} q') := \omega_\delta(q \xrightarrow{lab(e)/q} q') + 1$ 
11              $s := s \cdot q$ 
12              $q := q'$ 
13         case endElement do
14             let  $vp = s$  //  $p$  is top
15              $q' := ret(q, p, lab(e))$ 
16              $\omega_Q(q') := \omega_Q(q') + 1$ 
17              $\delta^{ret} := \delta^{ret} \cup \{q \xrightarrow{lab(e)/p} q'\}$ 
18              $\omega_\delta(q \xrightarrow{lab(e)/p} q') := \omega_\delta(q \xrightarrow{lab(e)/p} q') + 1$ 
19              $s := v$ 
20              $q := q'$ 
21         case characters do
22              $q' := int(q)$ 
23              $\omega_Q(q') := \omega_Q(q') + 1$ 
24              $\delta^{int} := \delta^{int} \cup \{q \xrightarrow{\tau} q' \mid \tau \in lab(e)\}$ 
25             for  $\tau \in lab(e)$  do  $\omega_\delta(q \xrightarrow{\tau} q') := \omega_\delta(q \xrightarrow{\tau} q') + 1$ 
26              $q := q'$ 
27   $F := F \cup \{q\}$ 
28   $\omega_F(q) := \omega_F(q) + 1$ 

```

---

of typing context  $u$  and left-sibling string  $v$ . Symbols  $\#$  and  $\$$  are a left-sibling separator and a placeholder for text.

- *Ancestor-based*. A state  $(u, v) \in (\Sigma^* \times (\Sigma \cup \{\$\}))^*$  is a pair of ancestor string and left-sibling string.
- *Ancestor-sibling-based*. A state  $(u, v) \in ((\Sigma \cup \{\$, \#\})^* \times (\Sigma \cup \{\$\}))^*$  is a pair of ancestor-sibling string and left-sibling string.

Initially, the intermediate VPA has a single nonaccepting start state  $(\epsilon, \epsilon)$ , no transitions, and counters are set to zero. Next states and transitions are created inductively from the start state, and counters are increased. For learning within the XSD language class, states must be ancestor based. Beyond XSD but within the 1PPT language class, states must be ancestor-sibling based.

3) *Local State Merging*: Based on Definition 10, every prefix of every datatyped event stream can characterize a



state. When complete ancestor, ancestor-sibling, and left-sibling strings are returned by state naming functions in Algorithm *incWeightedVPA*, the resulting automaton would accept exactly the learned documents similar to a PTA for regular languages. Generalization by state merging is then embedded in the state naming functions by returning equivalence classes of named states.

The distinguishing criterion is *locality*: two states are equal if they share the same  $l$ -local typing context and  $k$ -local left siblings. The refined naming functions are:

$$int_{k,l}(q) = (\pi_1(q), \sigma_k(\pi_2(q) \cdot \$)), \quad (4)$$

$$ret_{k,l}(q, p, e) = (\pi_1(p), \sigma_k(\pi_2(p) \cdot lab(e))), \quad (5)$$

$$call_{k,l}^{as}(q, e) = (\sigma_l(\pi_1(q) \cdot lab(e)), \epsilon), \quad (6)$$

$$call_{k,l}^{als}((r_1 \# \dots \# r_n, v), e) = (r_{n-l+1} \# \dots \# r_n \# \sigma_k(v \cdot lab(e)), \epsilon) \quad (7)$$

Ancestor- and ancestor-sibling-based naming schemes require a different *call* function denoted by superscripts *as* and *als* respectively. The suffix function  $\sigma_i(w)$  returns the  $i$ -length suffix of sequence  $w$ , and  $\pi_i(x)$  denotes the  $i$ th field of tuple  $x$ . For *characters* events,  $int_{k,l}$  is the same under both naming schemes; the typing context remains unchanged, and  $\$$  is appended to the left siblings as a placeholder. Using a placeholder for the next state is sound because of the mixed-content restrictions in Definition 6. For *endElement* events,  $ret_{k,l}$  is also the same under both naming schemes; the next state inherits the typing context from stack state  $p$ , and a new left sibling is added to the ones in  $p$ . In case of a *startElement* event, a new typing context is created, and left siblings are set to empty.

Parameters  $k$  and  $l$  specify the hypothesis space of the learner. For the lower bound  $k = l = 1$ , the learnable language class is a strict subclass of DTD. For  $l \geq 1, k = 1$ , both state naming schemes produce congruent automata, and the learnable language class is a strict subclass of XSD. Greater parameters increase the learnable language class, but also the state space grows, and more examples are necessary for convergence. If the true language class is not  $k$ - $l$ -local or when parameters are chosen too small, an approximation is learned.

4) *Generating a dXVPA*: The intermediate VPA and its counters still need to be translated into a dXVPA. Algorithm 4 (*trim*) creates a new intermediate VPA without zero-weight states and transitions. Furthermore, *trim* ensures a correct antichain of datatypes for internal datatype transitions between two states.

Algorithm 5 (*genXVPA*) generates a valid dXVPA from a trimmed intermediate VPA. States are partitioned into modules based on their typing context. The initial module  $m_0$  is the one called from state  $(\epsilon, \epsilon)$ . The *call* function from state naming guarantees that the entry of module  $m$  is always

---

**Algorithm 4: trim**


---

**Input:** VPA  $A = (Q, q_0, F, Q, \delta)$  over  $\Sigma \uplus T \uplus \bar{\Sigma}$   
lexical datatype system  $(T, \phi, \sim_s, \leq_s)$   
counters  $\omega_Q, \omega_F, \omega_\delta$

**Output:** VPA  $A' = (Q', q_0, F', Q', \delta')$

```

1  $\delta'^{call} := \delta^{call} \setminus \{q \xrightarrow{c/q} q' \mid \omega_\delta(q \xrightarrow{c/q} q') = 0\}$ 
2  $\delta'^{ret} := \delta^{ret} \setminus \{q \xrightarrow{\bar{c}/q} q' \mid \omega_\delta(q \xrightarrow{\bar{c}/q} q') = 0\}$ 
3  $\delta'^{int} := \delta^{int} \setminus \{q \xrightarrow{\tau} q' \mid \omega_\delta(q \xrightarrow{\tau} q') = 0\}$ 
4  $\delta''^{int} := \emptyset$ 
5 foreach  $\{(q, q') \mid \exists \tau. q \xrightarrow{\tau} q' \in \delta'^{int}\}$  do
6   let  $R = \{\tau \mid q \xrightarrow{\tau} q' \in \delta'^{int}\}$ 
7    $\delta''^{int} := \delta''^{int} \cup \{q \xrightarrow{\tau} q' \mid \tau \in \max_{\leq_{lex}} R\}$ 
8  $\delta' = \delta'^{call} \uplus \delta'^{ret} \uplus \delta''^{int}$ 
9  $Q' := Q \setminus \{q \mid \omega_Q(q) = 0\}$ 
10  $F' := F \setminus \{q \mid \omega_F(q) = 0\}$ 
```

---



---

**Algorithm 5: genXVPA**


---

**Input:** VPA  $A = (Q, q_0, F, Q, \delta)$  over  $\Sigma \uplus T \uplus \bar{\Sigma}$   
lexical datatype system  $(T, \phi, \sim_s, \leq_s)$   
**Output:** dXVPA  $A'$  over  $(\Sigma, M, \mu, T, \phi)$ , where  
 $A' = (\{Q_m, e_m, X_m, \delta_m\}_{m \in M}, m_0, X_{m_0})$

```

1  $M := \{u \mid (u, v) \in Q \text{ and } u \neq v\}$ 
2  $m_0 := u$  such that  $q_0 \xrightarrow{c/q_0} (u, \epsilon) \in \delta^{call}$ 
3 for  $m \in M$  do
4    $Q_m := \{(u, v) \in Q \mid u = m\}$ 
5    $e_m := (m, \epsilon)$ 
6    $X_m := \{q \in Q_m \mid q \xrightarrow{\bar{c}/p} q' \in \delta^{ret}\}$ 
7    $\delta_m^{call} := \{q \xrightarrow{c/q} q' \in \delta^{call} \mid q \in Q_m\}$ 
8    $\delta_m^{int} := \{q \xrightarrow{\tau} q' \in \delta^{int} \mid q, q' \in Q_m\}$ 
9    $\delta_m^{ret} := \{q \xrightarrow{\bar{c}/p} q' \in \delta^{ret} \mid q \in Q_m\}$ 
10   $\delta_m^{ret} := \delta_m^{ret} \cup \{q \xrightarrow{\bar{c}/p} q' \mid q \in X_m \text{ and } \exists q_m. q_m \xrightarrow{\bar{c}/p} q' \in \delta_m^{ret}\}$ 
11   $\delta_m = \delta_m^{call} \uplus \delta_m^{int} \uplus \delta_m^{ret}$ 
12  if  $\exists q. q \xrightarrow{c/q} e_m \in \delta^{call}$  then  $\mu(m) := c$ 
13  $A' := minimize(A')$ 
```

---

state  $(m, \epsilon)$ . Return transitions are added to all module exit states to ensure the single-exit property (Line 10).

Algorithm 6 (*minimize*) merges congruent modules. Kumar et al. [12] have shown that XVPA modules can be translated to DFAs, and this construction is extended to dXVPA modules. The algorithm compares modules  $m$  and  $n$ , and if they are reachable by the same element name and have congruent DFAs,  $n$  folds into  $m$  by redirecting calls and returns to corresponding states in  $m$ . The state bijection  $\varphi$  follows from bisimulation of the DFAs, and after a fold, *minimize* restarts until no fold occurs.

5) *Learner Properties*: Algorithm 7 assembles the learner. Incrementally updating an intermediate VPA prevents information loss from premature minimization of dXVPA modules. For a lexical datatype system, three nam-

---

**Algorithm 6: minimize**

---

**Input:** dXVPA  $A$  over  $(\Sigma, M, \mu, T, \phi)$ , where  
 $A = (\{Q_m, e_m, X_m, \delta_m\}_{m \in M}, m_0, X_{m_0})$

**Output:** minimized dXVPA  $A$

```
1 while  $\exists m \exists n. m, n \in M$  and  $m \neq n$  and  $\mu(m) =$   
    $\mu(n)$  and  $DFA_m \simeq DFA_n$  do  
2   let  $\varphi: Q_n \rightarrow Q_m$  // from bisimulation  
3   for  $q_n \xrightarrow{\bar{c}/p_i} q_i \in \delta_n^{ret}$  do  
4      $\delta_i^{call} := \delta_i^{call} \setminus \{p_i \xrightarrow{c/p_i} e_n\} \cup \{p_i \xrightarrow{c/p_i} e_m\}$   
5      $\delta_m^{ret} := \delta_m^{ret} \cup \{x_m \xrightarrow{\bar{c}/p_i} q_i \mid x_m \in X_m\}$   
6   for  $q_n \xrightarrow{\bar{c}/q_n} e_i \in \delta_n^{call}$  do  
7      $\delta_i^{ret} = \emptyset$   
8     for  $q_i \xrightarrow{\bar{c}/p_j} q_j \in \delta_i^{ret}$  do  
9       if  $j = n$  then  $\delta_i^{ret} := \delta_i^{ret} \cup \{q_i \xrightarrow{\bar{c}/\varphi(p_j)} \varphi(q_j)\}$   
10      else  $\delta_i^{ret} := \delta_i^{ret} \cup \{q_i \xrightarrow{\bar{c}/p_j} q_j\}$   
11     $\delta_i^{ret} := \delta_i^{ret}$   
12 if  $n = m_0$  then  $m_0 := m$   
13  $M := M \setminus \{n\}$  // remove module  $n$   
14  $\mu(n) := \emptyset$ 
```

---

ing functions, and parameters  $k$  and  $l$ , the incremental learner computes a dXVPA from document event stream  $w$ . The equivalent cXVPA can then check acceptance.

---

**Algorithm 7: Incremental learner**

---

**Input:** persistent VPA  $A$   
lexical datatype system  $dts = (T, \phi, \sim_s, \leq_s)$   
persistent counters  $\omega = (\omega_Q, \omega_F, \omega_\delta)$   
state naming  $f = (int_{k,l}, call_{k,l}, ret_{k,l})$  with  $k, l$   
document event stream  $w$

**Output:** dXVPA  $A'$

```
1 initially,  $A = (\{(\epsilon, \epsilon)\}, (\epsilon, \epsilon), \emptyset, \{(\epsilon, \epsilon)\}, \emptyset)$   
2  $A, \omega := incWeightedVPA(A, dts, f, \omega, w)$   
3  $A' := genXVPA(trim(A, dts, \omega))$ 
```

---

**Theorem 2.** The learner is (1) incremental, (2) set-driven, (3) consistent, (4) conservative, (5) strong-monotonic, and identifies a subclass of IPPT mixed-content XML.

Incremental learning follows from Algorithm 7. A set-driven learner follows from calling *incWeightedVPA* repeatedly for a set of examples and generating the dXVPA after the last one. Set-driven learning is insensitive to the order of presented examples, and this property follows from state naming and treating states and transitions as sets. A learner is consistent if all learned examples are accepted, conservative if a current hypothesis is kept as long as no contradicting evidence is presented, and strong-monotonic if the language increases with every example [35], [36]. These properties follow from updating sets of states and transitions in the intermediate VPA using the state naming functions.

States and call and return transitions are never deleted, and new ones are only added when observed. Also, an internal transition on datatype  $\tau$  is only removed if a new transition on  $\tau'$  is added, where  $\tau'$  covers  $\tau$ .

A learned dXVPA is always deterministic because of the restriction to  $k$ - $l$ -local IPPT. Checking acceptance using the corresponding cXVPA is therefore linear in the length of the document event stream.

### E. Anomaly Detection Refinements

The learning process could be targeted by poisoning [13], and two operations for dealing with adversarial training data are proposed.

---

**Algorithm 8: unlearn**

---

**Input:** VPA  $A = (Q, q_0, F, Q, \delta)$  over  $\Sigma \uplus T \uplus \bar{\Sigma}$   
lexical datatype system  $dts = (T, \phi, \sim_s, \leq_s)$   
counters  $\omega_Q, \omega_F, \omega_\delta$   
document event stream  $w$

**Output:** updated VPA  $A$  and counters  $\omega_Q, \omega_F, \omega_\delta$

```
1  $s := \perp$  // empty stack  
2  $q := q_0$  // current state  
3 for  $e$  in  $dtyped(w)$  do  
4   switch  $eventType(e)$  do  
5     case startElement do  
6        $q' := \delta^{call}(q, lab(e))$   
7        $\omega_Q(q') := \omega_Q(q') - 1$   
8        $\omega_\delta(q \xrightarrow{lab(e)/q} q') := \omega_\delta(q \xrightarrow{lab(e)/q} q') - 1$   
9        $s := s \cdot q$   
10       $q := q'$   
11     case endElement do  
12       let  $vp = s$  //  $p$  is top  
13        $q' := \delta^{ret}(q, lab(e), p)$   
14        $\omega_Q(q') := \omega_Q(q') - 1$   
15        $\omega_\delta(q \xrightarrow{lab(e)/p} q') := \omega_\delta(q \xrightarrow{lab(e)/p} q') - 1$   
16        $s := v$   
17        $q := q'$   
18     case characters do  
19        $q' := \delta^{int}(q, \tau)$  for some  $\tau \in lab(e)$   
20        $\omega_Q(q') := \omega_Q(q') - 1$   
21       for  $\tau \in lab(e)$  do  $\omega_\delta(q \xrightarrow{\tau} q') := \omega_\delta(q \xrightarrow{\tau} q') - 1$   
22        $q := q'$   
23  $\omega_F(q) := \omega_F(q) - 1$   
24  $A := trim(A, dts, \omega_Q, \omega_F, \omega_\delta)$ 
```

---

We distinguish poisoning attacks that are uncovered at some later time and poisoning attacks that remain hidden but are statistically rare. Therefore, *unlearning* removes a once learned example from the intermediate VPA, and *sanitization* trims low-frequent transitions and states.

Algorithm 8 (*unlearn*) simulates a run on the document event stream that needs to be forgotten, traverses the intermediate VPA, and decrements counters. The document must have been learned before at an earlier time for the operation to be sound.

---

**Algorithm 9:** *sanitize*

---

**Input:** VPA  $A = (Q, q_0, F, Q, \delta)$  over  $\Sigma \uplus T \uplus \bar{\Sigma}$   
lexical datatype system  $dts = (T, \phi, \sim_s, \leq_s)$   
counters  $\omega_Q, \omega_F, \omega_\delta$   
**Output:** updated VPA  $A'$  and counters  $\omega'_Q, \omega'_F, \omega'_\delta$

```
1 for any defined transition  $x$  do  $\omega'_\delta(x) := \omega_\delta(x) - 1$ 
2 for  $q \in Q$  do
3    $\omega'_Q(q) := \sum_{\text{transition } x \text{ to } q} \omega'_\delta(x)$ 
4   if  $q \in F$  then  $\omega'_F(q) := \omega'_Q(q)$ 
5  $A' := \text{trim}(A, dts, \omega'_Q, \omega'_F, \omega'_\delta)$ 
6 let  $Q_u$  be the unreachable states in  $A'$ 
7 if  $Q_u \neq \emptyset$  then
8   if  $(F' \setminus Q_u) = \emptyset$  then // revert changes
9      $\omega'_Q := \omega_Q; \omega'_F := \omega_F; \omega'_\delta := \omega_\delta; A' := A$ 
10  else // remove unreachable states
11    for  $q \in Q_u$  do
12      for any transition  $x$  to  $q$  do  $\omega'_\delta(x) := 0$ 
13       $\omega'_Q(q) := \omega'_F(q) := 0$ 
14     $A' := \text{trim}(A, dts, \omega'_Q, \omega'_F, \omega'_\delta)$ 
```

---

Algorithm 9 (*sanitize*) trims low frequent states and transitions by decrementing all counters. The algorithm has two stages. First, counters for all transitions are decremented, and counters of states are recomputed. Second, unreachable states are identified and decremented to zero for deletion. If no final state is reachable, all weight counters are restored because sanitization is not applicable.

It should be stressed that sanitization should only be applied after a large number of examples have been learned. The operation violates the consistent, conservative, and strong-monotonicity properties of the learner. Also, after a *sanitize* operation, *unlearn* becomes unsound.

### III. EXPERIMENTAL EVALUATION

The proposed approach has been implemented in Scala 2.11.7, and two aspects of performance are considered: detection performance and learning progress.

#### A. Measures

By assuming binary classification between *normal* and *attack*, the following performance measures are computed from labeled datasets: recall/detection rate ( $Re$ ), false-positive rate ( $FPR$ ), precision ( $Pr$ ), and  $F_1$  for overall performance [37]. Identification in the limit has a convergence point, but practical convergence can only be estimated by counting mind changes between incremental steps [38].

**Definition 11** (Mind changes). Mind changes  $MC_i$  are the number of states and transitions whose counters switched from zero to one after learning document event stream  $w_i$ .

Parameters  $k$  and  $l$  embody a strong combinatorial upper bound on the number of states and transitions for a finite

number of elements. In the worst case of randomness, convergence is reached when the state space is fully saturated.

#### B. Datasets

Table I summarizes the four datasets. The learner infers a dXVPA from training data, and performance is measured by validating the testing data with the corresponding cXVPA. Datasets Carsale and Catalog have been synthetically generated using the stochastic XML generator ToXGene [39]. For providing a realistic setting, a *VulnShopService* and a randomized *VulnShopClient* have been implemented for capturing SOAP messages. This Apache Axis2 1.6.0 SOAP/WS-\* web service uses Apache Rampart 1.6.0 for WS-Security and provides two service operations: regular shop orders (dataset VulnShopOrder) and digitally signed shop orders (dataset VulnShopAuthOrder). For realism, the implementation strictly followed the Axis2 and Rampart examples. The business logic utilized Java beans, and Java2WSDL automatically generated an Axis2 service from beans. Names for operations and Java classes have been deliberately chosen to require types in a schema.

Attacks in synthetic datasets were added manually. Attacks in the simulated datasets are recordings of actual attacks, e.g., WS-Attacker-1.7 [6] for Denial-of-Service (high node count, coercive parsing) and signature wrapping.

#### C. Performance

1) *Baseline Performance*: Schema validation using Apache Xerces 2.9.1 established a baseline, and results are listed in Table II. The schemas for the Carsale and Catalog datasets were extracted from ToXGene configurations, and simple types were set to datatype string or more informative datatypes when applicable. The VulnShopOrder and VulnShopAuthOrder datasets needed a schema collection from the web service because of the composed WS-\* standards.

The schemas in synthetic datasets are free from extension points, and schema validation achieved good performance as expected. The baseline for the simulated *VulnShopService* however illustrated the effect of extension points. Half of the attacks in VulnShopOrder were identified because of structural violations or datatype mismatches, but all Denial-of-Service attacks at extension points passed. Furthermore, no signature wrapping attack was identified.

2) *Detection Performance*: Table III summarizes the best results by the proposed algorithms for lowest parameters  $k$  and  $l$ . The best parameters were found in a grid search over values  $k, l \in \{1, \dots, 5\}$  and the two naming schemes.

The proposed language-based anomaly detection approach outperformed the baseline. No false positives were detected, and the best results were already achieved with the simplest parameters, i.e., ancestor-based state naming and  $k = l = 1$ . All structural anomalies caused by attacks were detected. It should be stressed that  $k = l = 1$  was a good-enough

Table I  
EVALUATION DATASETS

Dataset	Training Normal	Testing										
		Normal	Attack	XML tampering	High node count	Coercive parsing	Script injection	Command injection	SQL injection	SSRF attributes	XML injection	Signature wrapping
Carsale	50	1000	17	1	3	2	3	2	2	1	3	0
Catalog	100	2000	17	1	3	2	3	2	2	2	2	0
VulnShopOrder	200	2000	28	2	4	2	5	3	5	3	4	0
VulnShopAuthOrder	200	2000	78	0	0	0	0	0	0	0	0	78

Table II  
SCHEMA VALIDATION BASELINE PERFORMANCE

Dataset	$Pr$	$Re$	$FPR$	$F_1$
Carsale	100%	82.35%	0%	90.32%
Catalog	100%	76.47%	0%	86.67%
VulnShopOrder	100%	50%	0%	66.67%
VulnShopAuthOrder	undef.	0%	0%	undef.

Table III  
BEST PERFORMANCE USING ANCESTOR-BASED STATES

Dataset	$k$	$l$	$Pr$	$Re$	$FPR$	$F_1$
Carsale	1	1	100%	100%	0%	100%
Catalog	1	1	100%	82.35%	0%	90.32%
VulnShopOrder	1	1	100%	92.86%	0%	96.30%
VulnShopAuthOrder	1	1	100%	100%	0%	100%

approximation of the language to identify attacks, but more sound types were inferred for  $l > 1$ .

Some script and command injection attacks were not identified. These attacks have in common that exploitation code appears in texts and use CDATA fields to hide special characters, e.g., angled brackets and ampersands, from the XML parser’s lexical analysis. The lexical datatype system is too coarse in this case because the inferred datatype `normalizedString` permits the attack-identifying characters.

3) *Learning Progress*: Learning progress was measured in mind changes, and Figures 7a–7d summarize the fastest converging settings for the four datasets. When converged, the performance coincided with Table III. In every training iteration, the learner randomly drew a training document without replacement for learning, and the validator checked acceptance of testing data for measuring improvements. Because of randomness, runs were repeated 15 times, average values for  $F_1$  and  $FPR$  were computed, and the error regions in the plots illustrate minimal and maximal values in the random learning processes.

The first training example always caused many mind changes because there were no states and transitions yet. The strong-monotonicity property guarantees that detection performance either increases or stays the same after learning an example and assuming it is not a poisoning attack.

In the real world, detection performance is not observable but mind changes are. As shown in the figures, mind changes became less frequent over time, and a long period of zero mind changes could be a heuristic for convergence.

The quick convergence in Figure 7c and 7d stemmed from the simplicity of the language automatically generated by Java2WSDL. The generator only supports sequential (member variables) and iterating (arrays) productions but no choice. A few examples were sufficient for finding a good-enough approximation with small parameters  $k$  and  $l$ .

4) *Unlearning and Sanitization*: Unlearning reverses learning, and Figure 7e illustrates the effects. In this scenario, a successful attacker was able to feed poisoning attacks to the learner, and performance dropped accordingly. At a later time, a hypothetical expert identified the poisoning attacks and started unlearning them. The detection performance recovered, and knowledge gained in between attacks and unlearning remained in the model.

Sanitization trims low-frequent states and transitions. A single hidden poisoning attack was injected after 10% learning progress, and there was an impact on performance. After 75% progress, sanitization was performed. Figure 7f shows the effects of sanitization. In at least one of the 15 trials, the learner had no stable language representation at the moment of sanitization. Good knowledge was trimmed, performance dropped, and more mind changes after sanitization were necessary to recover again. Knowledge gained from a single example could be lost by sanitization.

#### IV. RELATED WORK

This work focuses on XML stream validation because of large documents and open-ended streams (e.g., XMPP). Stream validation has been introduced by Segoufin and Vianu [40] using finite-state machines and pushdown automata. Kumar et al. [12] consider document event streams as visibly pushdown languages (VPLs), a class of deterministic context-free languages, and the authors propose XVPAs as a better representation. XVPAs have therefore been extended with datatypes for text contents.

Schema inference from a set of documents focuses on finding simple regular expressions for schema productions. Beyond the expressiveness of DTD, Chidlovskii [41] and

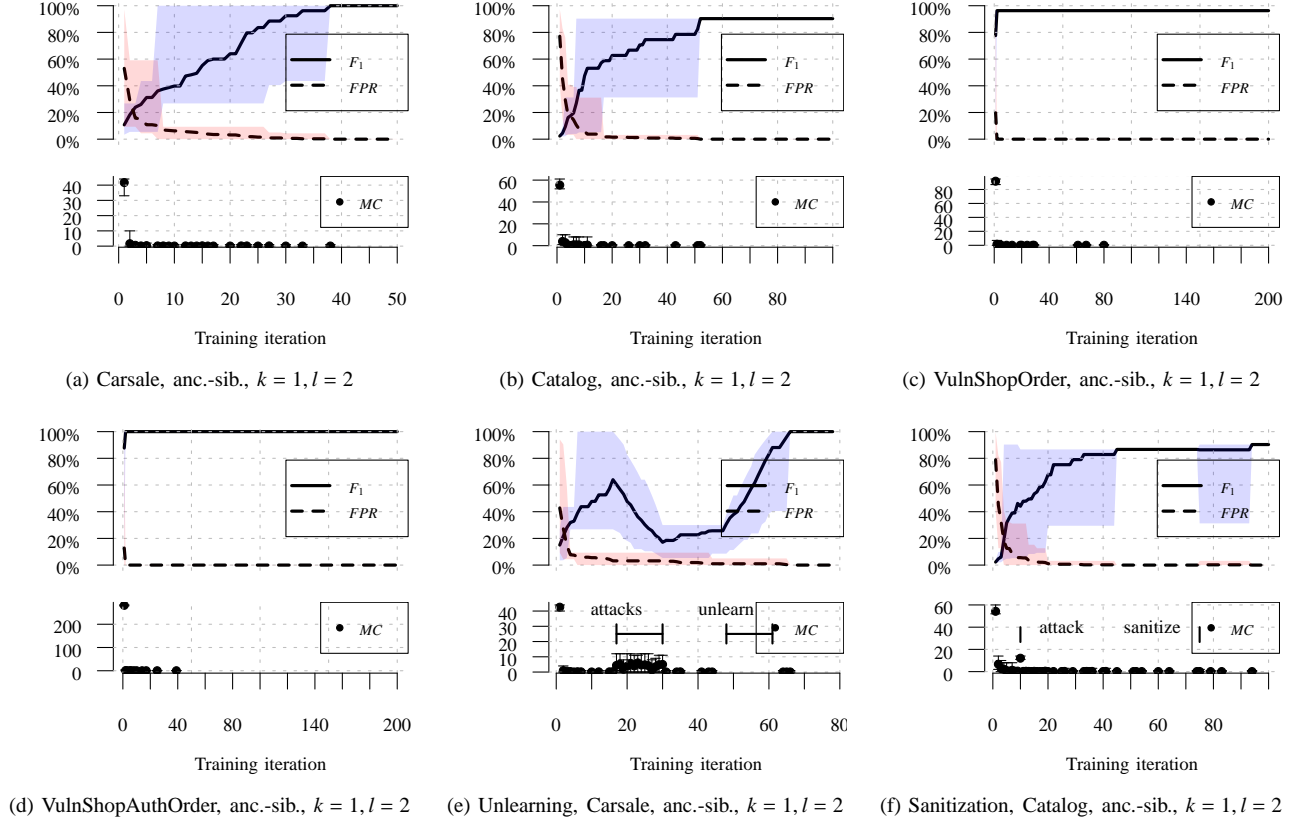


Figure 7. Learning progress

Mlýnková and Nečaský [42] propose grammar-based approaches, where info-set tree nodes turned into productions. These productions are then generalized by determinism constraints [41] and heuristics [42]. Bex et al. [43] propose schema inference in terms of tree automata, where up to  $k$  ancestor elements in a document characterize a type. This work has motivated the use of locality as a generalization strategy. Lexical subsumption for datatype inference was first mentioned by Chidlovskii [41] and Hegewald et al. [44]; however, not all XSD datatypes have been considered. The proposed approach considers a datatype choice instead of a single datatype, all distinguishable XSD datatypes are used, and a preference heuristic refines a choice.

With respect to anomaly detection, Menahem et al. [45] propose a feature extraction process for documents, so existing machine-learning algorithms can be reused, but structural information is lost. A schema is assumed to be available, and this direction has therefore not been further pursued. Another anomaly detection approach specifically for tree structures is based on geometry. Rieck [46] introduces tree kernels as measures of shared information between two parse trees. Kernels enable global and local anomaly detection, and this method could eventually be extended to XML info-set trees. Global anomaly detection finds a volume-minimal sphere that encloses the vector-embedded trees, and local

anomaly detection computes kernel-based distances to the nearest neighbors. Approximate tree kernels [47] are a trade-off for reducing computational costs. However, this method assumes a tree which conflicts with streaming requirements.

## V. CONCLUSIONS

This paper proposes a grammatical inference approach for learning the accepted language of an XML-based system. Schema validation is ineffective as a defense mechanism when extension points are present. For language-based anomaly detection, an automaton is inferred from examples, so documents with unexpected structure or text contents can be identified. It is also possible to translate such an automaton into a schema [12]. The contributions are dXVPAs as language representations for mixed-content XML, cXVPAs as an optimization of dXVPAs for efficient stream validation, algorithms for datatype inference from text, an incremental learner, and an experimental evaluation in synthetic and realistic scenarios.

The dXVPAs capture well-nested event streams, i.e., linearizations of trees, but no integrity constraints, to stay within a language class that allows efficient stream validation. This approach is nevertheless effective as a detection method because a learned language has no extension points. Improving the learning setting from  $k$ - $l$ -local languages

toward more powerful ones, e.g., by query learning [38], is a major open research question. Inferring and validating integrity constraints are also open research questions; however, Arenas et al. [48] have already shown that this problem is computationally much harder.

Simple parameters ( $k = 1, l = 2$ ) for the learner outperformed baseline schema validation in experiments; nonetheless, there are limitations. Some attacks in experiments could not be identified because lexical spaces of XSD datatypes are too coarse. Introducing more fine-grained datatypes would improve the detection rate. Also, repetitions are not bounded, and an order on unordered attributes is assumed. Repetition bounds and unordered attributes are two additional open research questions.

Finally, the unlearning and sanitization operations help to deal with adversarial training data, but the operations only apply after a poisoning attack has happened. The experiments indicated that the momentum of mind changes in the learning progress could be a heuristic for identifying a poisoning attack automatically while it is learned.

#### ACKNOWLEDGMENT

The research has been supported by the Christian Doppler Society, and the results were produced while the author was affiliated with the Christian Doppler Laboratory for Client-Centric Cloud Computing, JKU Linz, Austria.

#### REFERENCES

- [1] W3C, “Extensible Markup Language (XML) 1.0 (Fifth Edition),” Nov. 2008, accessed 2014-02-17. [Online]. Available: <http://www.w3.org/TR/2008/REC-xml-20081126/>
- [2] M. McIntosh and P. Austel, “XML signature element wrapping attacks and countermeasures,” in *Proc. SWS’05*. ACM, 2005, pp. 20–27.
- [3] W3C, “XML Schema Part 0: Primer Second Edition,” Oct. 2004, accessed 2014-07-14. [Online]. Available: <http://www.w3.org/TR/xmlschema-0/>
- [4] J. Somorovsky, M. Heiderich, M. Jensen, J. Schwenk, N. Gruschka, and L. Lo Iacono, “All your clouds are belong to us: Security analysis of cloud management interfaces,” in *Proc. CCSW’11*. ACM, 2011, pp. 3–14.
- [5] J. Somorovsky, A. Mayer, J. Schwenk, M. Kampmann, and M. Jensen, “On breaking SAML: Be whoever you want to be,” in *Proc. Security’12*. USENIX, 2012, pp. 21–21.
- [6] C. Mainka, J. Somorovsky, and J. Schwenk, “Penetration testing tool for web services security,” in *Proc. SERVICES’12*. IEEE, Jun. 2012, pp. 163–170.
- [7] S. Gajek, L. Liao, and J. Schwenk, “Breaking and fixing the inline approach,” in *Proc. SWS’07*. ACM, 2007, pp. 37–43.
- [8] S. Gajek, M. Jensen, L. Liao, and J. Schwenk, “Analysis of signature wrapping attacks and countermeasures,” in *Proc. ICWS’09*. IEEE, 2009, pp. 575–582.
- [9] M. Jensen, L. Liao, and J. Schwenk, “The curse of namespaces in the domain of XML signature,” in *Proc. SWS’09*. ACM, 2009, pp. 29–36.
- [10] M. Jensen, C. Meyer, J. Somorovsky, and J. Schwenk, “On the effectiveness of XML schema validation for countering XML signature wrapping attacks,” in *Proc. IWSSC’11*. IEEE, Sep. 2011, pp. 7–13.
- [11] H. Lampesberger, “A grammatical inference approach to language-based anomaly detection in XML,” in *Proc. ECTCM’13*. IEEE, 2013, pp. 685–693.
- [12] V. Kumar, P. Madhusudan, and M. Viswanathan, “Visibly pushdown automata for streaming XML,” in *Proc. WWW’07*. ACM, 2007, pp. 1053–1062.
- [13] G. F. Cretu, A. Stavrou, M. E. Locasto, S. J. Stolfo, and A. D. Keromytis, “Casting out demons: Sanitizing training data for anomaly sensors,” in *Proc. S&P’08*. IEEE, 2008, pp. 81–95.
- [14] W3C, “XML Information Set (Second Edition),” Feb. 2004, accessed 2014-08-02. [Online]. Available: <http://www.w3.org/TR/xml-infoset/>
- [15] Java, “JSR 173: Streaming API for XML,” Dec. 2013, accessed 2015-02-16. [Online]. Available: <https://www.jcp.org/en/jsr/detail?id=173>
- [16] OASIS, “RELAX NG,” Dec. 2001, accessed 2014-07-14. [Online]. Available: <https://www.oasis-open.org/committees/relax-ng/spec.html>
- [17] W. Martens, F. Neven, T. Schwentick, and G. J. Bex, “Expressiveness and complexity of XML Schema,” *ACM Trans. Database Syst.*, vol. 31, no. 3, pp. 770–813, 2006.
- [18] H. Wu, T. Ling, G. Dobbie, Z. Bao, and L. Xu, “Reducing graph matching to tree matching for xml queries with id references,” in *Proc. DEXA’10*, ser. LNCS. Springer Berlin Heidelberg, 2010, vol. 6262, pp. 391–406.
- [19] J. Lindsey, “Subtyping in W3C XML Schema, part 1,” *The Data Administration Newsletter*, Apr. 2008, accessed 2015-03-23. [Online]. Available: <http://www.tdan.com/view-articles/7185>
- [20] D. Stephenson, “XML Schema best practices,” HP, Tech. Rep., 2004, accessed 2015-03-25. [Online]. Available: <http://xml.coverpages.org/HP-StephensonSchemaBestPractices.pdf>
- [21] M. Jensen, N. Gruschka, and R. Herkenhöner, “A survey of attacks on web services,” *Computer Science - Research and Development*, vol. 24, no. 4, pp. 185–197, 2009.
- [22] A. Falkenberg, M. Jensen, and J. Schwenk, “Welcome to ws-attacks.org,” 2011, accessed 2015-02-05. [Online]. Available: <http://www.ws-attacks.org>
- [23] T. D. Morgan and O. A. Ibrahim, “Xml schema, dtd, and entity attacks,” Virtual Security Research, LLC, Tech. Rep., May 2014, accessed 2015-03-16. [Online]. Available: <http://www.vsecurity.com/download/papers/XMLDTDEntityAttacks.pdf>

- [24] W3C, “XML Signature Syntax and Processing (Second Edition),” Jun. 2008, accessed 2015-01-23. [Online]. Available: <http://www.w3.org/TR/xmlschema-0/>
- [25] N. Gruschka and L. Iacono, “Vulnerable cloud: SOAP message security validation revisited,” in *Proc. ICWS’09*. IEEE, Jul. 2009, pp. 625–631.
- [26] M. A. Rahaman, A. Schaad, and M. Rits, “Towards secure SOAP message exchange in a SOA,” in *Proc. SWS’06*. ACM, 2006, pp. 77–84.
- [27] E. M. Gold, “Language identification in the limit,” *Inform. Control*, vol. 10, no. 5, pp. 447–474, 1967.
- [28] H. Fernau, “Identification of function distinguishable languages,” *Theor. Comput. Sci.*, vol. 290, no. 3, pp. 1679–1711, 2003.
- [29] R. Alur and P. Madhusudan, “Visibly pushdown languages,” in *Proc. STOC’04*. ACM, 2004, pp. 202–211.
- [30] W3C, “W3C XML Schema Definition Language (XSD) 1.1 Part 2: Datatypes,” Apr. 2012, accessed 2015-04-02. [Online]. Available: <http://www.w3.org/TR/xmlschema11-2/>
- [31] G. J. Bex, F. Neven, and J. Van den Bussche, “DTDs versus XML Schema: A practical study,” in *Proc. WebDB’04*. ACM, 2004, pp. 79–84.
- [32] G. J. Bex, W. Gelade, F. Neven, and S. Vansummeren, “Learning deterministic regular expressions for the inference of schemas from XML data,” *ACM Trans. Web*, vol. 4, no. 4, pp. 1–32, 2010.
- [33] P. García and E. Vidal, “Inference of k-testable languages in the strict sense and application to syntactic pattern recognition,” *IEEE Trans. Pattern Anal.*, vol. 12, no. 9, pp. 920–925, 1990.
- [34] M. Murata, D. Lee, M. Mani, and K. Kawaguchi, “Taxonomy of XML schema languages using formal language theory,” *ACM Trans. Internet Techn.*, vol. 5, no. 4, pp. 660–704, 2005.
- [35] D. Angluin, “Inductive inference of formal languages from positive data,” *Inform. Control*, vol. 45, no. 2, pp. 117–135, 1980.
- [36] K. P. Jantke, “Monotonic and non-monotonic inductive inference,” *New Generat. Comput.*, vol. 8, no. 4, pp. 349–360, 1991.
- [37] J. Davis and M. Goadrich, “The relationship between Precision-Recall and ROC curves,” in *Proc. ICML’06*. ACM, 2006, pp. 233–240.
- [38] C. de la Higuera, *Grammatical Inference: Learning Automata and Grammars*. Cambridge University Press, 2010.
- [39] D. Barbosa, A. Mendelzon, J. Keenleyside, and K. Lyons, “ToXgene: A template-based data generator for XML,” in *Proc. SIGMOD’02*. ACM, 2002, pp. 616–616.
- [40] L. Segoufin and V. Vianu, “Validating streaming XML documents,” in *Proc. PODS’02*. ACM, 2002, pp. 53–64.
- [41] B. Chidlovskii, “Schema extraction from XML data,” Xerox, Tech. Rep. 2001/200, 2001, accessed 2015-06-20. [Online]. Available: <http://www.xrce.xerox.com/Research-Development/Publications/2001-200>
- [42] I. Mlýnková and M. Nečaský, “Heuristic methods for inference of XML schemas: Lessons learned and open issues,” *Informatica*, vol. 24, no. 4, pp. 577–602, 2013.
- [43] G. J. Bex, F. Neven, and S. Vansummeren, “Inferring XML schema definitions from XML data,” in *Proc. VLDB’07*. VLDB Endowment, 2007, pp. 998–1009.
- [44] J. Hegewald, F. Naumann, and M. Weis, “XStruct: Efficient schema extraction from multiple and large XML documents,” in *Proc. ICDEW’06*. IEEE, 2006, pp. 81–81.
- [45] E. Menahem, A. Schclar, L. Rokach, and Y. Elovici, “XML-AD: Detecting anomalous patterns in XML documents,” *Inform. Sciences*, 2015.
- [46] K. Rieck, “Machine learning for application-layer intrusion detection,” Ph.D. dissertation, TU Berlin, Germany, 2009.
- [47] K. Rieck, T. Krüger, U. Brefeld, and K.-R. Müller, “Approximate tree kernels,” *Journal Mach. Learn. Res.*, vol. 11, pp. 555–580, 2010.
- [48] M. Arenas, J. Daenen, F. Neven, M. Ugarte, J. V. D. Bussche, and S. Vansummeren, “Discovering XSD keys from XML data,” *ACM Trans. Database Syst.*, vol. 39, no. 4, pp. 28:1–28:49, Dec. 2014.